This is a preprint from arxiv.org:1404.6334
Final version has been published in Neural Computation, 27(5), pp 1102 – 1119, 2015 (May).
Please download the final version from the MIT Press Webpage
http://www.mitpressjournals.org/loi/neco

arXiv:1404.6334v5 [cs.NE] 8 Aug 2015

Input anticipating critical reservoirs show power law forgetting of unexpected input events

Norbert Michael Mayer[1]
[1]Department of Electrical Engineering and
Advanced Institute of Manufacturing with High-tech Innovations (AIM-HI),
National Chung Cheng University, Min-Hsiung, Chia-Yi, Taiwan
mikemayer@ccu.edu.tw

Keywords: Reservoir computing, scale invariance, power law

Abstract

Usually, reservoir computing shows an exponential memory decay. This paper investigates under which circumstances echo state networks can show a power law forgetting. That means traces of earlier events can be found in the reservoir for very long time spans. Such a setting requires critical connectivity exactly at the limit of what is permissible according the echo state condition. However, for general matrices the limit cannot be determined exactly from theory. In addition, the behavior of the network is strongly influenced by the input flow. Results are presented that use certain types of restricted recurrent connectivity and anticipation learning with regard to the input, where indeed power law forgetting can be achieved.

1 Introduction

Stability of a dynamical system is usually described in terms of the eigenvalues of the Jacobian, which clearly can identify a system as asymptotically stable or unstable. In the simplest scalar case one might consider an iterative system $x_{t+1} = f(x_t)$, with a fix-point at $x_\infty = f(x_\infty)$. As a standard approach (see e.g. [27]) one can analyze the asymptotic behavior by investigating a polynomial series of f around x_∞

$$f(x) \approx x_\infty + \alpha(x - x_\infty) + \beta(x - x_\infty)^2 + \gamma(x - x_\infty)^3 \dots \ . \tag{1}$$

Near the fix-point, one only has to consider the absolute value of α in order to predict the future of the dynamics as

$$x_t - x_\infty \propto \alpha^t. \tag{2}$$

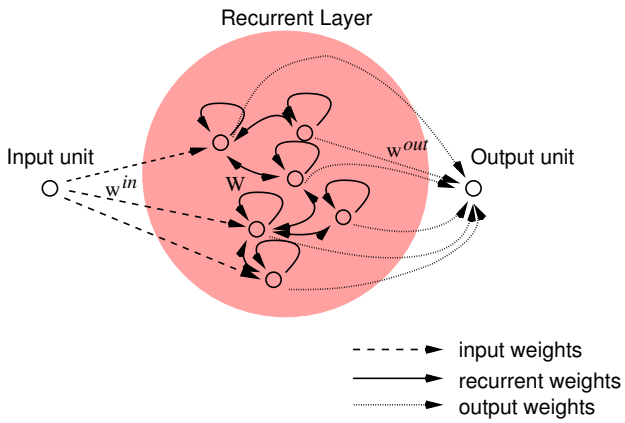

Figure 1: ESN scheme: The initial ESN approach chooses input and recurrent connectivity randomly, although it has to obey the echo state condition. Learning is only applied to the output layer. What is the *best* connectivity with regard to the other layers and some certain input statistics is still subject to an ongoing debate to which this work would like to contribute.

Thus, x_t converges to x_∞ if $|\alpha| < 1$, and it diverges from x_∞ if $|\alpha| > 1$. However, for the narrow class of systems of $|\alpha| = 1$ the linear term does not contribute to the convergence or divergence of x_t [1]. In this case the system can be considered critical in the sense that its dynamical properties can change if relatively small variations of α occur. Also, in the critical regime the dynamical system is controlled by the higher order terms (β, γ ...) of eq. 1. Among other consequences also the temporal convergence can then be approximated by a power law:

$$x_t - x_\infty \approx c\, t^b, \tag{3}$$

where c and b are constant values.

More general, power law properties in the dynamic variables and statistics can appear in critical states [7]. Vice versa, if power law statistics appear in experimental measurements, often a dynamical system in a critical regime can be assumed [1]. For example, in brain slices of rats, electrode measurements revealed such kind of statistics in spontaneous activity cascade lengths ('avalanches') [3]. If networks in the brain are dynamical systems in a critical regime this is an interesting feature which raises the question of the biological purpose of such a phenomenon.

In this paper a reservoir computing approach (for an overview over some recent developments see [15]) is presented where a critical dynamics is arranged around the expected input, that is if the network is trained with the expected input, it forgets any unexpected input in a power law fashion. Work of [23] can be seen as a hint that at

[1] This situation is also known as one of the cases where linear stability analysis fails (cf. [27], p. 93 ff.)

least with respect to untrained stimuli, some cortical areas can be interpreted as a kind of reservoir computing. In the scope of this work, I will introduce a neural networks approach that is based on Echo-State-Networks (ESNs). The ESN is a reservoir approach that uses rate-based neurons with real valued transfer functions, initially sigmoid functions. For my very basic considerations it turned out to be sufficient and can be easily tuned into a critical regime in the sense of this paper. The most important foundations of this paper's approach are those approaches that relate to critical ESNs [8, 4, 24, 6, 22].

The learning approach is specifically designed in a way that the network learns to anticipate the input, i.e. to predict it. This anticipation can also be labeled as balancing the network, a term which has been coined earlier [32] in the context of chaotic dynamics of integrate and fire neurons. The purpose of this mechanism is to prevent predictable information from entering the network, and producing activity in the recurrent layer. The intention is rather to let only the un-predicted input, i.e. special events with some level of unusualness, pass into the recurrent layer and stay there for further processing, while standard ESN approaches can only keep information about these events in the order of the half-time period of the forgetting factor.

In the following section I introduce the Echo-State paradigm and how it can be tuned into the critical state. I then discuss the role and implementation of the anticipation/balancing part of the algorithm. A result section follows. I conclude the work with a discussion concerning the potential information theoretic benefit of my model and biological implications.

2 Model

The model is based on Jäger's ESN approach [13, 14]. It consists of an input, recurrent layer and possibly an output layer (though not explicitly introduced in the scope of this paper). ESNs are composed of rate-based neurons with real valued transfer functions; the update rule is:

$$\begin{aligned} \mathbf{x}_{lin,t} &= \mathbf{W}\mathbf{x}_{t-1} + \mathbf{w}^{in}\mathbf{u}_t && (4) \\ \mathbf{x}_t &= \theta\left(\mathbf{x}_{lin,t}\right) && (5) \\ \mathbf{o}_t &= \mathbf{w}^{out}\mathbf{x}_t && (6) \end{aligned}$$

where the vectors $\mathbf{u}_t$,$\mathbf{x}_t$, and $\mathbf{o}_t$ are the input, the neurons of the hidden layer, and the neurons of the output layer, respectively, and $\mathbf{w}^{in}$, $\mathbf{W}$, and $\mathbf{w}^{out}$ are the matrices of the respective synaptic weight factors.
As a convention the transfer function $\theta(.)$ is continuous, differentiable and monotonically increasing with the limit $1 \geq \theta'(.) \geq 0$, which is compatible with the requirement that $\theta(.)$ fulfills the Lipschitz continuity with $L = 1$.
Jäger's approach uses random matrices for $\mathbf{W}$ and $\mathbf{w}^{in}$, learning is restricted to the output layer $\mathbf{w}^{out}$ (see fig. 1). The learning (i.e. training $\mathbf{o}_t$) can be performed by linear regression. Since the learning process with regard to some output itself is not of interest in the scope of this paper, this part of the approach is not outlined here.

2.1 Echo-state condition

A necessary condition for the performance of an ESN network is that the echo state condition is fulfilled. Consider a time-discrete recursive function $\mathbf{x}_{t+1} = F(\mathbf{x}_t, \mathbf{u}_t)$, where the $\mathbf{x}_t$s are interpreted as internal states and the $\mathbf{u}_t$s form some external input sequence, i.e. the stimulus. The definition of the echo-state condition is the following. Assume an infinite stimulus sequence: $\bar{\mathbf{u}}^\infty = \mathbf{u}_0, \mathbf{u}_1, \ldots$ and two random initial internal states of the system $\mathbf{x}_0$ and $\mathbf{y}_0$. To both initial states $\mathbf{x}_0$ and $\mathbf{y}_0$, the sequences $\bar{\mathbf{x}}^\infty = \mathbf{x}_0, \mathbf{x}_1, \ldots$ and $\bar{\mathbf{y}}^\infty = \mathbf{y}_0, \mathbf{y}_1, \ldots$ can be assigned;

$$\begin{aligned} \mathbf{x}_{t+1} &= F(\mathbf{x}_t, \mathbf{u}_t) \\ \mathbf{y}_{t+1} &= F(\mathbf{y}_t, \mathbf{u}_t). \end{aligned} \tag{7}$$

Then the system $F(.)$ is called *universally state contracting* if **independent** from the set $\mathbf{u}_t$ and for **any** ($\mathbf{x}_0$,$\mathbf{y}_0$) and all real values $\epsilon > 0$, there exists an iteration τ for which

$$d(\mathbf{x}_t, \mathbf{y}_t) \leq \epsilon \tag{8}$$

for all $t \geq \tau$.
Jäger showed ([13] pp. 43) that the echo state condition is fulfilled if and only if the network is universally state contracting. The ESN is designed to be universally state contracting and thus to fulfill the echo-state condition. Heuristics show that the performance of ESNs becomes better near the critical point where the echo state condition is just narrowly fulfilled. The rest of this paper is dedicated to determining under which circumstances the time series of eq. 8 can converge as

$$d(\mathbf{x}_t, \mathbf{y}_t) \precsim t^b, \tag{9}$$

rather than the usual

$$d(\mathbf{x}_t, \mathbf{y}_t) \precsim a^t, \tag{10}$$

which connects reservoirs to the considerations from the introduction and eq. 1, i.e. to critical connectivity.
Critical connectivity can be achieved by just narrowly fulfilling Jäger's conditions on the recurrent weight matrix of a network that has echo states:
A necessary condition is:

- C1 A network has echo states only if the absolute value of the biggest eigenvalue of $\mathbf{W}$ is below 1.

A sufficient condition is:

- C2 A network has echo states if the biggest singular value of $\mathbf{W}$ is smaller than one[2].

[2] A closer sufficient condition has been found in [5, 34]. It is slightly better than the one outlined here, but still leaves a gap to the necessary condition for general matrices.

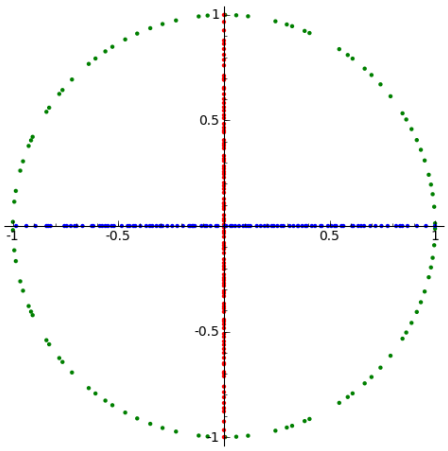

Figure 2: Plot of eigenvalues of samples of different types of normal matrices are depicted according to their positions on the complex plane (x represents the real part, y the imaginary part), the eigenvalues of a symmetric matrix are blue, eigenvalues of a skew-symmetric matrix are red and eigenvalues of an orthogonal matrix are green.

2.2 Normal recurrent connectivity matrices W

A general ESN becomes critical – for some input sequences – somewhere in the range between C1 and C2. In the case of **normal** matrices for the recurrent connectivity $\mathbf{W}$, i.e. matrices that fulfill:

$$\mathbf{W}^T\mathbf{W} = \mathbf{W}\mathbf{W}^T, \tag{11}$$

both conditions are either true or false at the same time. The proof can be found for example in [11].

Several prominent types of matrices are normal: symmetric, orthogonal($\mathcal{O}(n)$), permutation and skew symmetric.

These different types have advantages and disadvantages with respect to learning algorithms: First of all permutation matrices can be discarded: Although untrained reservoirs of permutation matrices tend to show a good performance compared to other types of connectivity patterns, they have a significant disadvantage with regard to optimization tasks. For a recurrent layer with n neurons there are $n!$ different permutations possible, which forms a discrete set of disjunct points rather than a continuous manifold within the parameter space. So they seem hardly usable here since fine tuning cannot be performed by continuous parameter optimization. Symmetric matrices are less useful since such reservoirs show only trivial dynamics in real valued reservoirs, i.e. the eigenvectors are mapped to themselves. Orthogonal matrices are a good choice. They form a group with respect to matrix multiplication, their absolute eigenvalues are all 1 (see fig. 2); the complete set of all orthogonal matrices forms a manifold that can be parameterized. Special virtues of orthogonal matrices have already been investigated [33] in a different context. Skew-symmetric also might be useful. This may be biologically plausible as a result of STDP [16]. However, considering fig. 2 one can see that the eigenvalues are all pure imaginary and more or less equally distributed along the imaginary axis. That is why – similar to symmetric matrices – most components of the signal

encoded in the hidden neurons decay exponentially and only the information along the eigenvectors belonging to the two largest eigenvalues ($\pm i$) survive for a power law time scale. Thus, since usually only two eigenvectors have the highest absolute eigenvalue, the information carried in the critical state might be poor. In the scope of this work, orthogonal matrices are being investigated. Previously also skew-symmetric matrices have been checked. Technically, one can obtain a skew-symmetric or orthogonal matrix by adding additional steps into the learning process that – in every iteration – enforce the respective constraints.
In the present case one has to enforce either one of the two following constraints

- C3 skew-symmetric $\mathbf{W}$ or
- C4 orthogonal $\mathbf{W}$.

In the case of skew-symmetric matrices, one also has to set

- C5 largest absolute eigenvalue/singular value is set to 1.

Orthogonal matrices already have all their absolute eigenvalues set to 1. Thus, if the network has to be tuned exactly into the critical state one way to do that is to either enforce C3 and C5 or C4 in every iteration. For practical purposes it suffices in many cases to use general matrices with a largest absolute eigenvalue set to one (i.e. only enforce condition C5).

2.3 Transfer function and anticipation

In order to identify critical behavior, one has to come back to eq. 8, which basically defines convergence between two different state sequences if the input is identical and the network connectivity is also identical. The most direct way to enforce the convergence is to design the network as contractive, i.e. to ensure in each iteration t:

$$d(\mathbf{x}_{t+1}, \mathbf{y}_{t+1}) = l(t) \cdot d(\mathbf{x}_t, \mathbf{y}_t), \tag{12}$$

where $l(t)$ indicates whether the distance between $\mathbf{x}_t$ and $\mathbf{y}_t$ converges to zero. If $l(t)$ becomes and stays equal or larger than 1 from any t, the two sequences do not converge anymore and the network does not have the echo state property. On the other hand if for all t and an upper limit L one has the inequality

$$l(t) \leq L < 1, \tag{13}$$

then network is an ESN.
Eq. 12 defines $l(t)$ as

$$l(t) = \frac{d(\mathbf{y}_{t+1}, \mathbf{x}_{t+1})}{d(\mathbf{y}_t, \mathbf{x}_t)} = \frac{||\theta(\mathbf{W}\mathbf{y}_t + I) - \theta(\mathbf{W}\mathbf{x}_t + I)||_2}{||\mathbf{y}_t - \mathbf{x}_t||_2}, \tag{14}$$

where $I = \mathbf{w}^{in}\mathbf{u}_t$ is the influx from the input to the network, and $||.|\,|_2$ is the Euclidean norm. For tiny differences between $\mathbf{x}_t$ and $\mathbf{y}_t$ one can approximate $l(t)$ as

$$l(t) \approx \frac{||\dot{\Theta}\mathbf{W}(\mathbf{y}_t - \mathbf{x}_t)||_2}{||\mathbf{y}_t - \mathbf{x}_t||_2} \leq ||\mathbf{J}||_2, \tag{15}$$

where $\dot{\Theta}$ is a diagonal matrix with entries $\dot{\theta}(\mathbf{W}x + I)$, and the right hand expression is the spectral norm of the Jacobian of the neural network dynamics, i.e.

$$J_{ij} = \dot{\theta}(x_{lin,i})W_{ij} \tag{16}$$

where $x_{lin,i}$ is the ith component of the vector $\mathbf{x}_{lin,t}$ from eq. 4. Note that by definition of the spectral norm, $l(t)$ approximates $||\mathbf{J}||_2$ if the difference between $\mathbf{y}_t$ and $\mathbf{x}_t$ is co-linear to the largest eigenvector of $\mathbf{J}$.
With respect to criticality the choice of the transfer function is important. On one hand note that the ESN requires convergence for **any** input sequence which explains why $1 \geq \dot{\theta}$. On the other hand, if one can achieve

$$1 = \dot{\theta}(x_{lin,i}), \tag{17}$$

one gets simply $\mathbf{J} = \mathbf{W}$, and thus $||\mathbf{J}||_\mathbf{2} = \mathbf{1}$. This case is analogous to the situation of eq. 1 and the case $|\alpha| = 1$ there; the linear terms do not affect the dynamics anymore. Instead, the dynamics are rather controlled by higher order terms of Taylor series around the critical point. The shape of the higher terms depends on the general shape of the transfer function; see [18] on an analysis of critical state, in particular under which circumstances the echo state condition can be preserved.
Then one fundamental problem of creating critical behavior is to design the network in a way that the $x_{lin,i}$s are always tuned to values where eq. 17 is fulfilled. The way to this is to let the network anticipate, i.e. predict the next input. Ideas that lead to the prediction of input have been outlined in earlier works [20, 25, 31].

Usually $\theta = \mathsf{tanh}(.)$ is used, and it is used in the initial approach. Note that the critical point of eq. 17 is reached at

$$\mathbf{x}_{lin,t} = [0, 0, 0 \dots], \tag{18}$$

i.e. there is only one possible critical state. This is un favorable as one can see from the following considerations.
The optimization of the hidden layer is going to redirect the linear response resulting from the expected input exactly to the vector above. Since in the case of the hyperbolic tangent there is only one possible value, i.e. all linear responses are 0.
On the one hand, the optimization process will result in the same single response of the entire network for all expected inputs. On the other hand, the point of the recurrent network is to resolve ambiguities by accounting for the input history. This is not possible if the total network is indifferent to what the previous input was if this input was expected.
So there is no information transfer from the previous state to the next state. For the current purpose, it makes sense to use the following transfer function:

$$\theta(x) = 0.5x - 0.25\,\mathsf{sin}(2x) \tag{19}$$

It has to be noted that the maximal derivative $\theta'(x)$ is 1 at $x_{lin,i} = \pi(n+1/2)$, where n is an integer number. Usually for small input values, only the lowest states $x_{lin,i} = \pm\pi/2$ are used by the network, resulting in a setting resulting in a setting where the critical point of eq. 17 is reached as

$$\mathbf{x}_{lin,t} = [\pm\frac{\pi}{2}, \pm\frac{\pi}{2}, \pm\frac{\pi}{2} \dots]. \tag{20}$$

This results in 2^N different possible states for the total network, where N is the number of the hidden neurons. Thus, information transfer within the critical state is possible.

2.4 Learning anticipation in the hidden layer

Several learning rules for the hidden layer have been proposed. In many cases an information theoretic measure has been applied [24, 29]. In the present approach it is intended to follow the idea to anticipate the input, which also has been tried in different ways previously [8].

I propose here the minimization of the cost-function

$$E(\mathbf{W}, \mathbf{w}^{in}) = \sum_i < cos(\mathbf{x}_{lin,t,i})^2 >_t, \tag{21}$$

where $\mathbf{x}_{lin,t,i}$ is the i^{th} component of $\mathbf{x}_{lin,t}$ (see eq. 4). The cost function becomes minimal if $\mathbf{x}_{lin,t,i} = \pi(n + 1/2)$, which fits to the critical points of the transfer function eq. 19. Note that $\mathbf{x}_{lin,t}$ also contains input to the system at the time t. Thus, the optimization includes an 'anticipation'of the input, i.e. some input is expected and already counted in order to come close to $\pm\pi/2$ for each neuron. The optimization is done on both $\mathbf{W}$ and $\mathbf{w}^{in}$ by gradient descent

$$\begin{aligned} \Delta\mathbf{W} &= -\nabla_{\mathbf{W}} E \\ \Delta\mathbf{w}^{in} &= -\nabla_{\mathbf{w}^{in}} E, \end{aligned} \tag{22}$$

where E is the cost-function of eq. 21, i.e. one gets

$$\begin{aligned} \Delta\mathbf{W}_{ij} &= -2\,\mathsf{cos}(\mathbf{x}_{lin,t,i})\,\mathsf{sin}(\mathbf{x}_{lin,t,i})\,\mathbf{x}_{t-1,j} \\ \Delta\mathbf{w}^{in}_{ij} &= -2\,\mathsf{cos}(\mathbf{x}_{lin,t,i})\,\mathsf{sin}(\mathbf{x}_{lin,t,i})\,\mathbf{u}_{t,j.} \end{aligned} \tag{23}$$

The update is done by

$$\begin{aligned} \tilde{\mathbf{W}} &= \mathbf{W}_{\mathsf{old}} + \epsilon\Delta\mathbf{W} \\ \mathbf{w}^{in}_{\mathsf{new}} &= \mathbf{w}^{in}_{\mathsf{old}} + \epsilon\Delta\mathbf{w}^{in}, \end{aligned} \tag{24}$$

where $\mathbf{W}_{\mathsf{old}}$ and $\mathbf{w}^{in}_{\mathsf{old}}$ are recurrent and input connectivity matrices of the previous iteration, ϵ represents the learning rate, and $\mathbf{w}^{in}_{\mathsf{new}}$ is the updated input connectivity matrix. As a result of the learning process the updated $\tilde{\mathbf{W}}$ is not necessarily critical, respectively normal in the sense of eq. 11 anymore. Thus, it remains to find a normal matrix that is as near as possible to the result of the learning step. Thus, in each iteration after adaptation the conditions C4 and C5 can be enforced by setting

$$\begin{aligned} [\mathbf{U}, \mathbf{S}, \mathbf{V}^T] &= \mathsf{SVD}(\tilde{\mathbf{W}}) \\ \mathbf{W}_{\mathsf{new}} &= \mathbf{U} \cdot \mathbf{V}^T, \end{aligned} \tag{25}$$

where SVD represents the singular value decomposition routine of standard mathematical toolboxes, usually the return value is a list composed of three matrices, i.e.

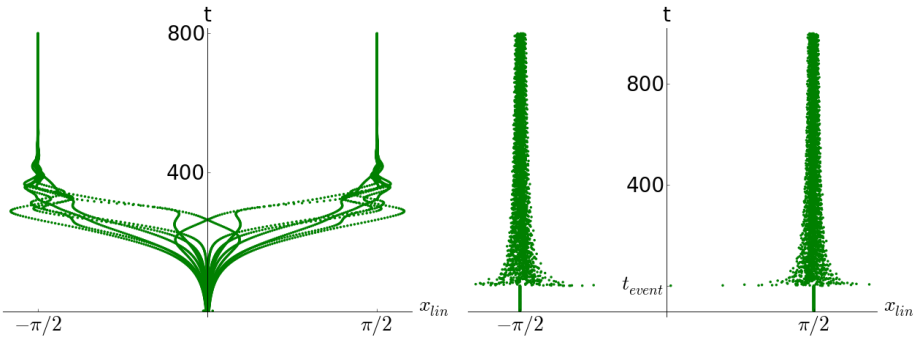

Figure 3: Linear response of the 2 state test setting. Left: Linear response during the learning process. Right: Linear response of network after an un-expected input event. One can see slow convergence process to the target values $\pm\pi/2$.

$[\mathbf{U}, \mathbf{S}, \mathbf{V}^T]$ where $\mathbf{U}$ and $\mathbf{V}$ are two orthogonal matrices and $\mathbf{S}$ is a positive semi-definite diagonal matrix. Since the constraint is enforced in every iteration, already $\mathbf{W}_{\text{old}}$ is an orthogonal matrix. After the learning step, if we assume ϵ to be small, one can expect $\mathbf{S}$ to be near an identity matrix. Setting $\mathbf{W}_{\text{new}}$ according to eq. 25 makes it an orthogonal matrix since the right hand side of eq. 25 is a product of 2 orthogonal matrices. So the recurrent connectivity can forced back on the manifold of orthogonal matrices. In this way condition C4 is enforced. Since all absolute eigenvalues of orthogonal matrices are equal to one, the condition C5 is also enforced by eq. 25. For symmetric matrices as a different type of normal matrices other ways of enforcing constraints can be applied For example, as outlined in [17], C3 and C4 can be enforced by

$$\mathbf{W}_{new} = \frac{\tilde{\mathbf{W}} - \tilde{\mathbf{W}}^T}{\max \text{abs}(\text{eigv}(\tilde{\mathbf{W}} - \tilde{\mathbf{W}}^T))} \tag{26}$$

where eigval represents a function that computes numerically a list of eigenvalues of a matrix. In the case of general matrices C5 can approximately enforced by

$$\mathbf{W}_{new} = \frac{\tilde{\mathbf{W}}}{\max \text{abs}(\text{eigv}(\tilde{\mathbf{W}}))}. \tag{27}$$

3 Simulation results

3.1 Reduced model

The first experiment was done with a model of 8 neurons and a normalized orthogonal connectivity matrix in the recurrent layer. The input was one neuron alternating between 1 and -1. In this case the learning algorithm for the hidden layer converges rapidly ($\approx$ 5000 iterations with a learning rate of 0.01). Fig. 3 depicts on the left side the linear response of each neuron ($\mathbf{x}_{lin,t}$) during the learning process. One can see that the learning follows the intended effect, i.e. the linear responses become more and more

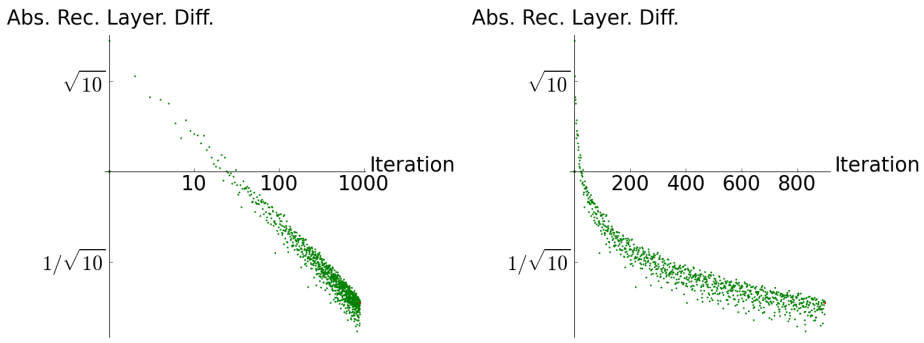

Figure 4: Left: Double log plot of the reduced model. The x-axis depicts time (iterations), and the y-axis depicts the difference measure between the undisturbed network with a network that received one unexpected input at one iteration. One can see that the decay is proportional to t^a, which results in a straight line of points in the double log plot. Iterations are counted from the moment at which one network receives one unexpected input. Right: Log-lin plot of the same data.

accurately either $-\pi/2$ or $\pi/2$. Note that with different initializations other values of the set $\pi/2 + n \times \pi$ can be approximated, where n is an integer number.
As a second step 2 identical copies of the network are created after the learning process. Then an unexpected input ($\mathbf{u} = -1$ instead of $\mathbf{u} = 1$) is presented to one of these networks at one time step. Fig. 3, right side, depicts the recording of the linear response of each of the 8 neurons, from briefly before the unexpected input until several hundred iterations afterwards. One can see that, although the neurons converge to either $-\pi/2$ or $\pi/2$, this process is very slow.
In order to get a quantitative picture, one can also take a metrical measure, in an analogous fashion to eq. 8[3].
Results are depicted in fig. 4. The results indicate that the reduced model shows the echo state property, however the decay is not exponential but a power law. Therefore, for a very long time period traces of the single unexpected event can be found in the network. Experiments using skew-symmetric matrices have been conducted previously and are going to be published in [17]

3.2 Stochastic sequences with temporal inference

A similar input model as in [21] was chosen to set up a subsequent test. Here an input sequence composed of 5 different state types labeled $\mathsf{A} - \mathsf{E}$ (see fig. 5) was used. Subsets of 4 subsequent states always belong together. The subset is either ABAD or ACAE, both occur i.i.d. randomly with equal probability. The sequence implies a simple grammar, i.e. the occurrence of a D in the state sequence requires the occurrence of a B two iterations before, while the occurrence of a E requires a C. Thus the network should infer the occurrence of the state D from the previous occurrence of the state B and the

[3] The applied metric was $d = \sum_i \mathsf{abs}(x_i - y_i)$, where x and y are the respective states of both networks.

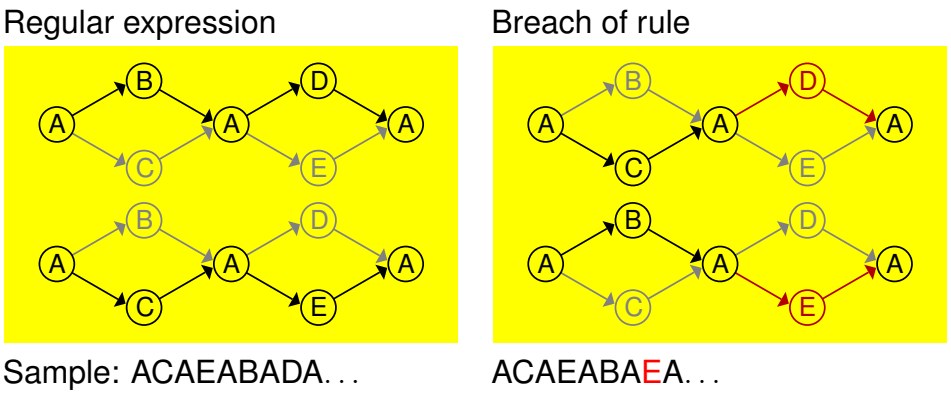

Figure 5: Grammar of the training set: Left the trained sequence of states. A regular sequence of four steps consists of either the sequence ABAD or ACAE. The occurrence of both sequences is random with equal probability. After four steps the sequence is repeated from the beginning again. For testing the networks, 2 identical networks were used where one network at one step received input at one time from a sequence that violates the implicitly trained rule.

occurrence of E from the previous occurrence of C.
Each of the states are encoded into different input vectors. In the present case a 2 dimensional input vector was implemented to encode each of the five states:

$$\begin{array}{lcl} \mathsf{A} & \rightarrow & \begin{pmatrix} 0 & 0 \end{pmatrix} \\ \mathsf{B} & \rightarrow & \begin{pmatrix} -1 & 0 \end{pmatrix} \\ \mathsf{C} & \rightarrow & \begin{pmatrix} 1 & 0 \end{pmatrix} \\ \mathsf{D} & \rightarrow & \begin{pmatrix} 0 & -1 \end{pmatrix} \\ \mathsf{E} & \rightarrow & \begin{pmatrix} 0 & 1 \end{pmatrix}. \end{array}$$

For this experiment general matrices were used and a largest eigenvalue of 1 was enforced. A network of 15 neurons was capable of virtually reaching the theoretical limit of the cost function during training. Instead of training an output function, the present approach again restricts itself to presenting a single unpredicted input and comparing the time series of internal states to an initially identical network where except for the one step over the whole time series, the same input is presented to the network. During the training phase, the network runs 20000 iterations. Initially the recurrent orthogonal matrix is scaled to $S_{max} = 0.8$. During the first 7500 iterations S_{max} is exponentially increased to 1.0. The learning rate of the gradient descent is set to 0.009. The learning process usually reaches values of the cost function that are around or smaller than 10^{-20}. After a transient of 1000 iterations, an exact copy of the first ESN is created. After that three different tests are performed.
In the first test, one network perceives a grammatical error at one iteration (see fig. 5, right side), i.e. both get the same time series as stimuli, except for a single occurrence at which the first networks receives an E according to the grammar and the other network receives an F, which violates the implicit rule of the training set. Left side of fig. 6 depicts how the difference between the 2 internal state vectors develops over time. The first iteration is identical to the iteration where the different inputs occur.

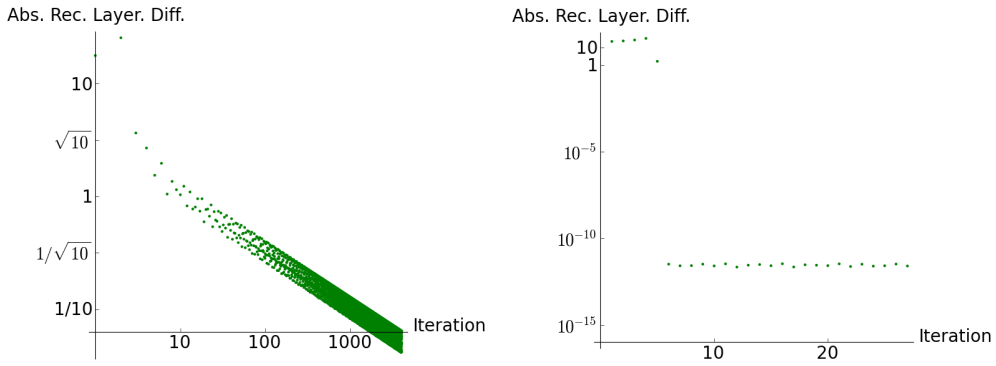

Figure 6: The plots show the convergence between two identical networks where one network receives a different input from the other network for a brief time. Left: Here one network receives at one iteration an input that represents a grammatical error in the context of the training set whereas the other network receives the grammatically correct sequence. The plot is depicted in a double logarithmic way, thus revealing power law descent of the metric distance between the two networks. The exponent of the decay is approximately -0.5. Right: Semilogarithmic plot of two networks that receive different, however in both cases grammatically correct input, at 2 time steps. Except for iteration 1 and iteration 3 both networks receive the same input. Thus, the first network perceives the sequence ABAD, whereas the second network gets ACAE.

As one can see the single violation results in an power law convergence. The metric distance between the two networks is

$$||\mathbf{x}_t - \mathbf{y}_t|| \propto t^{-b}, \tag{28}$$

where b is approximately $\frac{1}{2}$.
In contrast to the first result, if both networks receive 2 different input events that do not represent a violation of the grammar, the convergence of both networks is reached after 5 iterations (see fig. 6, right side). So in this case the trained network dynamics can return to default after one cycle of training patterns.
Finally, the third experiment tests what happens if both networks receive permanent violations in every iteration after receiving one iteration of differing input. Results are shown in fig. 7. As one can see the convergence is exponential in this case. One way to understand this result is to consider the information capacity of the reservoir as limited. As the network is not trained to predict the input, the unusual input drives the network further from the critical state.

4 Discussion

This paper is intended to contribute to the up-coming discussion about neural networks with critical connectivity [28]. Results presented here may contribute on both technical aspects of the study of biological neural networks such as the human cortex.

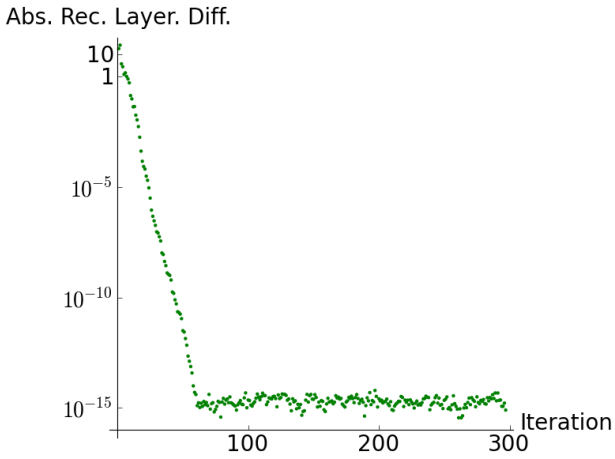

Figure 7: In this plot initially one network is exposed to violations of the grammatical rule, whereas the other network still receives the correct input. After that both networks receive permanent input with the wrong grammar. The plot is log-lin, which demonstrates that the difference between both networks vanishes exponentially in this case.

Technically, a system has been presented in which memories fade slower than exponential. In the present approach, that target could be achieved by applying two components into reservoir computing, i.e.

- **critical connectivity in the recurrent layer**,
- **an input prediction system**, a system that anticipates the next input in a way that it redirects the expected activity exactly into the critical point of the system.

These two components may also be important features in future approaches of reservoir computing and input driven systems in general.
With regard to the technical aspects of critical connectivity, note no heuristic measure to find the critical point has been applied.This differs from previous approaches that investigated the behavior reservoirs near the edge of chaos [8, 4, 24, 6, 22] . The graphs presented in the results section show a power law convergence over several orders of magnitude, which is only possible if the critical point is hit exactly. On one hand this is presumably not achievable by using the methods applied in earlier approaches. On the other hand, this is necessary to produce the plots that depict power law forgetting over several orders of magnitude. Further research is necessary how the restriction to normal matrices instead of the general ones affects the dynamics in the reservoir in a negative way.
Further, prediction of the input is essential for the results of this paper. Anticipation in reservoirs is possibly an important aspect for the future of ESNs. In addition to the considerations that are outlined in sect. 2.3, another information theoretic argument supports this notion. The idea is here that the total information content of the network is limited. In a non-predictive model the constant information influx requires that knowledge from previous experiences vanishes in an exponential fashion. No matter how the memory is organized, this is inevitable because the capacity of the reservoir is limited.

The predictive/anticipative approach uses a primitive type of memory compression which allows a single unexpected event to stay longer than exponential in the memory. In addition the present work could also be seen as an optimal predicting machine in the sense of Still's definition [30].
Related research can be found in the field of heterogeneous neural models [12, 10], where single input events can also be remembered over long time periods. In the case of LSTM, this can be achieved by adding specific gates in each unit that protect certain information from being forgotten.
Moreover, there could also be biological implications. Although the human memory system is by far a more complex mechanism than those that are outlined in this work, potential implications could be drawn from the present results. Different from machines human memory works often in a way that humans can remember unexpected events better. We usually can neither recall what we ate for lunch 78 days ago, nor precisely reconstruct the way we brushed our teeth last Monday. Events that are irrelevant or predictable from previous information or both are filtered away.
Also, human memory does not work from the first day, an effect known as childhood amnesia [9]. A little child has early experiences that are not stored as repeatable events in later life. Rather it is plausible to assume that these early childhood experiences form a framework of expectations of usual daily events. Only the deviations from these expectations are stored and stay as memories at later stages of life.
Another important hint that a mechanism similar to what is outlined in this paper is at work in the cortex are the ubiquitous predictive properties of the cortex recently have become common knowledge in the community [2].

Calculations for this manuscript are published online

The numerical calculations have been performed by using the open source sagemath.org [26] package. A part of the calculations are available under [19].

5 Acknowledgements

The Ministry of Science and Technology (MOST) of Taiwan provided the budget for our project - project numbers : 103-2221-E-194 -039 and 102-2221-E-194 -050. Also thanks go to AIM-HI for various ways of support. I also thank Chris Shane for his cross reading.

References

[1] Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality: An explanation of 1/f noise. *Physical Review Letters*, 59(4):381–384, 1987.

[2] Moshe Bar. *Predictions in the brain: Using Our Past to Generate a Future*. Oxford University Press, 2011. ISBN 978-0195395518.

[3] J. Beggs and D. Plenz. Neuronal avalanches in neocortical curcuits. *J. Neurosci.*, 24(22):5216–5229, 2004.

[4] Joschka Boedecker, Oliver Obst, Joseph Lizier, N. M. Mayer, and Minoru Asada. Information processing in echo state networks at the edge of chaos. *Theory in Biosciences*, 131:205–213, 2012.

[5] Michael Buehner and Peter Young. A tighter bound for the echo state property. *IEEE Transaction on Neural Networks*, 17(3):820–824, 2006.

[6] Lars Büsing, Benjamin Schrauwen, and Robert Legenstein. Connectivity, dynamics, and memory in reservoir computing with binary and analog neurons. *Neural Computation*, 22:1272–1311, 2010.

[7] Walter Greiner, Ludwig Neise, and Horst Stöcker. *Thermodynamik und Statistische Mechanik*. Verlag Harri Deutsch, 1993. Sect. 17.2: Kritische Indizes, pp. 512.

[8] Márton Albert Hajnal and András Lörincz. Critical echo state networks. In *S. Kollias et al. (Eds.), ICANN 2006, Part I, LNCS 4131*, pages 658–667, 2006.

[9] J. S. Herbert and O. Pascalis. *Memory development*. Oxford University Press, 2007.

[10] S Hochreiter and J Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.

[11] Roger A. Horn and Charles R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, 1991. pp. 157.

[12] Gregor M. Hörzer, Robert Legenstein, and Wolfgang Maass. Emergence of complex computational structures from chaotic neural networks through reward-modulated hebbian learning. *Cerebral Cortex*, 24(3):677–690, 2014.

[13] Herbert Jäger. The echo state approach to analysing and training recurrent neural networks with an erratum note. In *GMD Report 148, GMD German National Research Insitute for Information Processing*, 2010 2001.

[14] Herbert Jäger. Adaptive nonlinear system identification with echo state networks. In *Advances in Neural Information Processing Systems, NIPS 2002*, pages 593–600. MIT Press, Cambridge, MA, 2003 2002.

[15] Herbert Jäger, Wolfgang Maass, and Jose Principe (Eds.). Echo state networks and liquid state networks (special issue). *Neural Networks*, 20:287–432, 2007.

[16] Henry Markram, Wulfram Gerstner, and Per Jesper Sjöström. Spike-timing-dependent plasticity: a comprehensive overview. In *Frontiers in synaptic neuroscience*. Frontiers Media SA, Lausanne, 2012.

[17] N. M. Mayer. Adaptive critical reservoirs with algebraic forgetting of unexpected input sequences. In *S. Wermter et al. (Eds.), Artificial Neural Networks and Machine Learning ICANN 2014, LNCS 8681*, pages 49 – 56. 2014.

[18] N. M. Mayer. Proof about weakly universal state contracting networks. arXiv preprint arXiv:1411.6757, URL, `http://arxiv.org/pdf/1411.6757.pdf`, 2014.

[19] N. M. Mayer. Python sagemath code to produce figs. 3-7. URL, `http://crudescientists.org:8000/home/pub/34/` & `http://crudescientists.org:8000/home/pub/35/`, 2014.

[20] N. M. Mayer and Matthew Browne. Self-prediction in echo state networks. In *Proceedings of The First International Workshop on Biological Inspired Approaches to Advanced Information Technology (BioAdIT2004), Lausanne*, pages 40 – 48, 2004.

[21] N. M. Mayer, Oliver Obst, and Yu-Cheng Chang. Time series causality inference using echo state networks. In *Vincent Vigneron et al. (Eds.), Latent Variable Analysis and Signal Separation, LNCS 6365*, pages 279 – 286. Springer, 2010.

[22] T. Natschläger, N. Bertschinger, and R. Legenstein. At the edge of chaos: Real-time computations and self-organized criticality in recurrent neural networks. In *L. K. Saul et al. (Eds.), Advances in Neural Information Processing Systems, NIPS 17*, pages 145 – 152, 2005.

[23] Danko Nikolić, Stefan Häusler, Wolf Singer, and Wolfgang Maass. Distributed fading memory for stimulus properties in the primary visual cortex. *PLoS Biol*, 7(12):e1000260, 2009.

[24] Oliver Obst and Joschka Boedecker. Guided self-organization of input-driven recurrent neural networks. In *M. Prokopenko (Ed.), Guided Self-Organization: Inception*, pages 319–340. Springer, 2014.

[25] René Felix Reinhart and Jochen Jakob Steil. Regularization and stability in reservoir networks with output feedback. *Neurocomputing*, 90:96–105, 2012.

[26] Sagemath. URL, `http://www.sagemath.org`, 2014.

[27] Edward R. Scheinermann. *Invitation to Dynamical Systems*. Prentice Hall and Dower Publishers, 1996, 2012. Sect. 3.1, pp. 79 ff.

[28] HG Schuster, D Plenz, and E Niebur. *Criticality in Neural Systems*. John Wiley & Sons, 2014.

[29] Jochen Jakob Steil. Online stability of backpropagation-decorrelation recurrent learning. *Neurocomputing*, 69(79):642–650, 2006.

[30] Susanne Still, David A. Sivak, Anthony J. Bell, and Gavin E. Crooks. Thermodynamics of prediction. *Phys. Rev. Lett.*, 109:120604, 2012.

[31] David Sussillo and L.F. Abbott. Transferring learning from external to internal weights in echo-state networks with sparse connectivity. *PLOS One*, 2012. DOI: 10.1371/journal.pone.0037372.

[32] C. van Vreeswick and H. Sompolinsky. Chaotic balanced state in a model of cortical circuits. *Neural Computation*, 10:1321–1371, 1998.

[33] Olivia L. White, Daniel D. Lee, and Haim Sompolinsky. Short-term memory in orthogonal neural networks. *Phys. Rev. Lett.*, 92:148102, 2004.

[34] I. B. Yildiz, H. Jäger, and S. J. Kiebel. Re-visiting the echo state property. *Neural Networks*, 35:1–20, 2012.